Common Knowledge, Ambiguity, and the Value of Information in Games

Hailin Liu
Department of Philosophy
Carnegie Mellon University, Pittsburgh
hailinl@andrew.cmu.edu

Abstract

This paper asks whether the salient result about non-negative value of cost-free information holds in the context of games. By reexamining Osborne's example where information may hurt, it argues that the failure of this result is mainly driven by the assumption of common knowledge in the traditional framework of incomplete information games, since it leads to act-state dependence in a sequential setting. This paper also shows that such a failure occurs when we extend the framework of incomplete information games to allow for a representation of uncertainty using sets of probabilities and the use of Γ-maximin. Nevertheless, the key to this negative result is that a phenomenon called dilation of sets of probabilities obtains in this generalized setting.

Keywords. Bayesian games, value of information, ambiguity, act-state dependence, dilation, Γ-maximin.

1 Introduction

It is well known that in an individual decision problem a Bayesian decision maker should not refuse to receive cost-free information. To understand this result intuitively, note that choosing the expected utility maximizer d^* in the absence of new information has the same expectation as the plan of always making decision d^* no matter what additional information is forthcoming. But it might be that the decision d^* does not always maximize expected utility upon on receiving new information. It follows that the expected utility of choosing d^* initially cannot be greater than the prior expectation of choosing the best decisions after having more information. This ensures that additional information cannot be harmful to one's prior expectations. Thus it is rational for a Bayesian decision maker to postpone her terminal decision in order to acquire cost-free information. Indeed, such an intuitive idea can be traced back to Ramsey (1990), and has been formalized by Good (1967). However, this satisfactory result about the non-negative value of cost-free information fails to hold in many cases. For instance, Kadane et al. (2008) have shown that certain modifications of standard expected utility theory may require a decision maker to strictly prefer less information to more, thereby implying a negative value of information to the individual.

For our purposes, two cases under which this positive result does not obtain are worth mentioning. The first involves the idea of *act-state dependence*, namely probabilistic dependence between act and state, which is precluded by conventional expected utility theory. For example, within a small geographic market, consumers' inquiry about the price of a certain good may cause its price to rise. In this case, because of act-state dependence, a potential consumer strictly prefers not to learn the (cost-free) information about the price. As a result, the value of information is negative to the consumer. Another relevant instance of such a violation occurs in the context of decision making under uncertainty, where uncertainty is assumed to be represented by sets of probabilities rather than a single probability distribution. When extending expected utility theory to accommodate uncertainty aversion, it may happen that the set of unconditional probabilities for an event is properly included in the set of probabilities conditional on every event of some partition, which is a phenomenon known as *dilation* in the literature[1]. Given the presence of dilation, it should come as no surprise that a rational decision maker may refuse a free offer to learn a piece of new information. Therefore, the result introduced at the beginning of the paper is not robust with respect to the introduction of act-state dependence, as well as the choice of the modeling of uncertainty in the setting of single-agent decision making.

Moving beyond individual decision making, it has also been demonstrated that, in the context of games, more information may hurt the player who possesses the information[2]. More specifically, the player may be worse off when she has more information than when she does not. The simple logic behind this negative result is the following. In a game where one player has certain information

[1] See Seidenfeld and Wasserman (1993), Herron et al. (1994), and Herron et al. (1997) for an extensive and systematic study of the phenomenon of dilation.

[2] See for instance Akerlof (1970) and Osborne (2004) for several prominent examples that illustrate this observation concerning the negative value of information in games.

about the situation, this fact being commonly known among the players has a strategic impact on the players' rational strategy choices, which results in an inferior equilibrium outcome than the one for a game with more information. Thus, if that specific player can first choose between these two games and then play the chosen one, the optimal strategy is to play the game with less information, implying that new information has a negative value for that player. In light of this, we can conclude that the familiar result about the non-negative expected worth of cost-free information is not robust with respect to the introduction of strategic interaction either.

These findings suggest a need for a reconsideration of the information value problem in games. The aim of this paper is to provide a better understanding of this issue in games based on previously known results in the literature of decision theory. As a starting point of our investigation, we review the example introduced by Osborne (2004) exhibiting the counterintuitive result that in the context of Bayesian games a rational player may strictly prefer less information to more. In addition, the example helps us clarify the central issue and highlight the sequential problem involved in the comparison between a game where one player has less information, and a different version of the game where the player has more information. It is within such a sequential setting that the issue of whether information has a positive value can be properly assessed.

Next we show that the non-negative value of information can be restored by a weakening of the common knowledge assumption. Our analysis of Osborne's example makes it clear that the assumption of common knowledge plays a crucial role in the result of the negative information value in games. This leads us to consider a variant of the original game in which one player has more information while the other players are not aware of this fact. By comparing it with the original game, we show that more information to a player does lead to a better equilibrium outcome and thus has a positive value for the player[3]. We can intuitively understand this result by recognizing that act-state dependence is the real factor that has driven the result of the negative information value in games. The lack of common knowledge merely provides a convenient way to abstract from the kind of probabilistic dependence between a player's choice of the games and her probability about the opponent's choices.

We then deal with the question of whether the finding about the negative information value in games still holds in the presence of ambiguity. First, we briefly discuss how the traditional framework of Bayesian games developed by Harsanyi (1967/68) can be naturally extended to accommodate the idea of employing sets of probabilities to model uncertainty. Such an extension allows the players' initial beliefs about the state to be represented by closed and convex sets of probability distributions, which is more normatively appropriate and empirically grounded than the modeling of uncertainty through a single precise prior[4]. Following the pioneer work of Kajii and Ui (2005), we introduce a solution concept called Γ-*maximin equilibrium* that generalizes the concept of Bayesian Nash equilibrium to incomplete information games under ambiguity. Roughly speaking, our solution concept requires that each player chooses the optimal strategy in the sense of maximizing her minimum expected utility for each realization of her private signals. Using this solution concept, we then show that the phenomenon of the negative value of information reappears in the context of incomplete information games under ambiguity. This finding is perhaps not that surprising given the discussion on Osborne's example.

Nevertheless, we further argue that there exists an important distinction between Osborne's observation and the finding presented here. Specifically, we demonstrate that, when ambiguity is present, the conclusion about the negative value of information in games is robust with respect to the relaxation of the common knowledge assumption. This stands in a sharp contrast to what happens in Bayesian games when the assumption of common knowledge is weakened. In order to account for this difference, we briefly describe the phenomenon of dilation of sets of probabilities, and suggest that the role dilation together with the use of Γ-maximin plays in our examples is fundamental to understanding the result about the negative value of information in games under ambiguity. This implies that our finding bears a closer relationship to the results reported in Seidenfeld (2004) and Kadane et al. (2008) concerning the effect of dilation on the value of information in the context of single-agent decision making.

The remainder of this paper is organized as follows. In Section 2 we examine the well-known result about the negative value of information in Bayesian games in detail, and also present a different approach for determining the value of information in games by relaxing the assumption of common knowledge. Section 3 shows by examples that a non-expected utility player may pay not to receive cost-free information in incomplete information games under ambiguity. Section 4 discusses whether our finding still holds true even if there is a lack of common knowledge in games, and then relates it to the phenomenon of dilation. Section 5 contains a few concluding remarks.

[3] Neyman (1991) similarly argues that a player cannot be worse off by having more information provided that other players are not aware of it. The author thanks an anonymous reviewer for pointing out this reference.

[4] See for instance Knight (1921), Ellsberg (1961), Levi (1974), and Walley (1991) for a number of compelling arguments that justify the idea of using imprecise probabilities to model uncertainty.

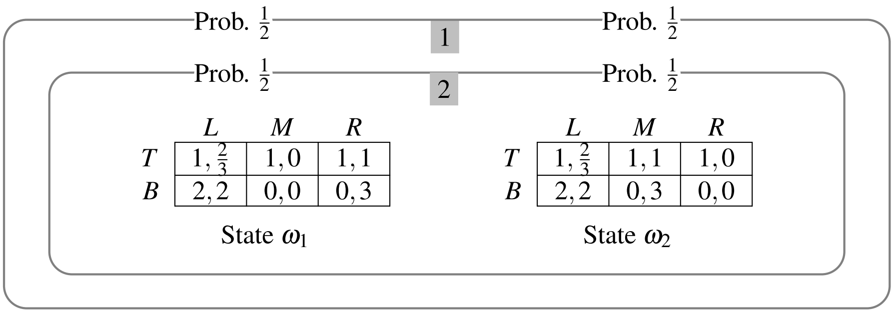

Figure 1: A Bayesian game with no information about the states

2 Common Knowledge and Negative Value of Information in Games

In this section we reconsider the example presented by Osborne (2004) where a player may assign a negative value to cost-free information in the context of Bayesian games. That is, the player becomes worse off in the version of the game where she has more information. It thus follows that the player would rationally pay to avoid learning this piece of new information, even though there is no cost associated with it. To keep things simple, we will return to this example later and argue that the assumption of common knowledge plays a critical role in Osborne's analysis, which throws considerable light on the nature of this negative result in games.

Example 2.1. *Consider the game structure shown in Figure 1. In this game, there are two players and two possible states $\{\omega_1, \omega_2\}$. The action set of player 1 is given by $A_1 = \{T,B\}$ and that of player 2 is given by $A_2 = \{L,M,R\}$. Players' payoffs under each strategy profiles are specified by two numbers in the corresponding box of the following table, with the first number being the payoff of player* 1. *Moreover, we assume that both players do not know the state and assign probability $\frac{1}{2}$ to each state. That is, both players believe that each state will occur with probability $\frac{1}{2}$, although neither of them knows which state they are actually in. And no further information concerning the state will be revealed to the players. Following normal convention, we further assume that everything about the game is common knowledge.*

Game theorists commonly regard the notion of Bayesian Nash equilibrium as the reasonable solution concept for solving Bayesian games. Informally speaking, a *Bayesian Nash equilibrium* consists of a collection of strategies such that each player's strategy is a *best reply* to the strategies the other players have chosen. It is easy to see by expectation that player 2's strategy L strictly dominates the other strategies. And player 1 has a unique best reply against L, namely, the action B. Hence, the strategy profile (B,L) is the unique Bayesian Nash equilibrium for the Bayesian game considered in this example. More importantly, note that this equilibrium gives rise to the outcome $(2,2)$ with the first component being the payoff to player 1.

It is well known in the decision-theoretic literature that in single-agent decision problems it is weakly better for an expected utility maximizer to wait for more information prior to make a terminal decision. In other words, every expected utility decision maker prefers more information to less. In order to test whether this is valid within the context of games, Osborne suggests to contrast Example 2.1 described above with the following case.

Example 2.2. *As before, we assume that both players do not know the state before receiving their private information and assign probability $\frac{1}{2}$ to each state. However, player 2 learns the state from her private information, whereas player 1 does not. In other words, after having received her private information, player 2 knows whether she is playing the strategic game on the left or the one on the right. By contrast, player 1 still does not know the state in this interim stage and thinks that with equal probability $\frac{1}{2}$ she is playing one of these two games. Nevertheless, player 1 knows the fact that player 2 is informed of the state. In a similar way, this strategic situation can be described by Figure 2 where the two frames labeled 2 enclosing each table indicate that player 2 can perfectly distinguish between these two tables.*

Notice that each type of player 2 has a strictly dominant action, namely, R and M respectively. This means that player 2 should choose to play the strategy (R,M), no matter what player 1 intends to do. Knowing that player 2 is informed of the state, player 1 can then anticipate the above reasoning on the behalf of player 2. Given this, player 1 should respond optimally by playing T, which is the unique best response to (R,M). Therefore, this game has a unique Bayesian Nash equilibrium $(T,(R,M))$, which gives rise to the outcome $(1,(1,1))$.

Now suppose that in a sequential setting player 2 is first asked to choose either to play the Bayesian game in Example 2.1, or to play the Bayesian game in Example 2.2, and then enters the chosen game. Note that player 2 would obtain more information about the state in Example 2.2 than

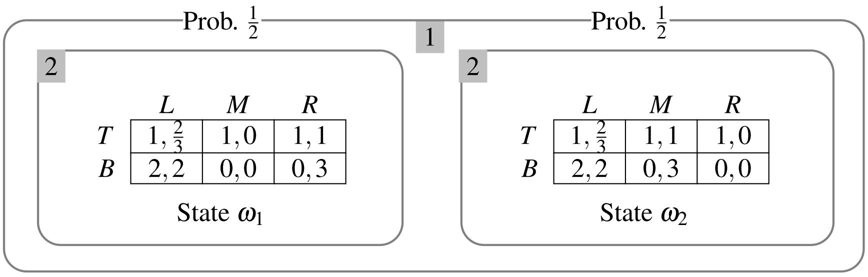

Figure 2: A Bayesian game in which player 2 knows the state

she would in Example 2.1. Contrary to our intuition, player 2 would rather choose to play the game in Example 2.1 instead of the one in Example 2.2, where she is actually informed of the state. It is important to remark, however, that such a choice of player 2 is perfectly rational in the exact sense of maximizing expected utility, since the unique Bayesian Nash equilibrium of the game in Example 2.1 yields player 2 a higher expected payoff than the unique Bayesian Nash equilibrium of the game in Example 2.2 does. This implies that player 2 strictly prefers to play the game with less information rather than more, which stands in stark contrast to the familiar result in decision theory about the non-negative value of information.

Generally, more information is valuable for making better decisions. In a strategic situation, however, knowing only that a certain kind of information becomes available to someone may alter one's behavior, even if she herself does not know what the information is. It is thus important to keep in mind that one should separate the real issue concerning the value of information in games, from the question whether the knowledge of others holding certain private information plays a role in determining its value. Next we suggest an instructive way to investigate whether more information is better for a player by restricting others from knowing that fact. Such an approach then avoids the kind of complication just described.

Example 2.3. *Consider a slightly modified version of the Bayesian game introduced in Example 2.2. As before, we assume that both players do not know the state before receiving their private information and assign probability $\frac{1}{2}$ to each state. Upon receipt of their private information, player 2 learns the state, whereas player 1 still does not know the state. Unlike in the previous game, in this case we assume that player 1 does not know the fact that player 2 is informed of the state. To be a bit more specific, player 1 still believes that player 2 assigns probability $\frac{1}{2}$ to each state like she does. Given that player 1 is not informed of player 2's updated belief about the state, it is obvious that the conventional assumption about common knowledge is no longer valid in the current case.*

The only difference between Example 2.2 and Example 2.3 lies in the fact that in the latter case we specifically make the information concerning player 2 knowing the state unavailable to player 1. Nevertheless, such a slight modification has greatly changed the strategic interaction between these two players. To see this, note that in Example 2.3 two players have rather different information about the situation: At the interim stage where players learn of their private information, player 1 does not know player 2's actual belief about the state, whereas player 2 has a complete knowledge of the strategic situation, including her own information being unknown to player 1. In this sense, player 1 holds a false belief about player 2's updated belief. In this setting, the usual strategic impact caused by new information goes away, since the information would only affect the person who has it. By restricting player 1 from knowing the fact that player 2 learns the state, we can then ask whether this piece of new information has any value to player 2 or not.

We first need to figure out how to resolve these kinds of strategic situations involving "imperfect" beliefs, as the notion of Bayesian Nash equilibrium is designed to solve only standard Bayesian games. As a minimal requirement of rationality, a reasonable solution should respect the information available to each player. Moreover, we want to follow the tradition of modeling the players as expected utility maximizers. Given these requirements, it seems reasonable to suggest that each player should choose a strategy that maximizes her expected payoff given the (possibly false) beliefs about the state and the strategies chosen by the other players, as long as these beliefs can be justified in terms of information available to her.

Now let us apply the above idea to the situation described in Example 2.3. First, we claim that player 1's optimal choice is to play B. To see this, note that in the interim stage player 1 still believes that player 2 does not know the state and assigns probability $\frac{1}{2}$ to each state. Given such a belief, player 1 would believe that player 2 will choose the action L, which strictly dominates the other two actions. Anticipating this conjecture, player 1 responds optimally by playing B. On the other hand, player 2 has perfect knowledge about the situation, even including the fact that player 1 holds incorrect belief about whether she knows the state. Since player 2 learns the state at the interim stage, player 2 will choose to play her strictly dominant actions in each state, that is, R and M respectively. Although player 2 can anticipate player 1 to select B rather than her equilibrium strategy L, player 2 would keep her optimal choice of (R,M) unchanged, since it still maximizes her expected payoff given B. Hence, the recommended play in this case is the strategy choice $(B,(R,M))$.

It is important to note that player 2's payoff under the optimal play in Example 2.3 is 3, which is greater than her payoff of 2 associated with the equilibrium of the game in Example 2.1. Thus, if player 2 is faced with an initial choice between whether to play the game in Example 2.1 or instead to play the game in Example 2.3, the rational decision is to choose the latter one. It follows that player 2 strictly prefers to have more information rather than less, which is in accordance with the familiar result concerning the non-negative value of cost-free information. This contrasts sharply with the foregoing analysis.

Note that common knowledge plays a critical role in solving both games in terms of Bayesian Nash equilibrium. In both games, everything about the games is common knowledge, and thus both players can reason about each other's strategy choice on the behalf of her opponent. In Example 2.1, based on her prior belief about the state, player 2 select her best action L, which leads player 1 to choose B. Similarly, in Example 2.2, player 2 utilizes her information of the state and singles out the optimal strategy (R,M), which induces player 1 to choose T. So, when player 2 is faced with the sequential problem of choosing first between these two games and then playing the chosen one, there exists probabilistic dependence between her choice of the games and her probability about the opponent's strategy choice. In this sense, act-state dependence arises in such a sequential problem.

On the other hand, act-state dependence does not arise when player 2 is asked to choose first whether to play the game in Example 2.1 or to play the game in Example 2.3, and then to play the selected game. To see this, recall that in Example 2.3 player 1 does not know the fact that player 2 has more information. In the light of this assumption, player 1 is not able to arrive at the same conclusion as player 2 does. Instead, player 1 would obtain the same conjecture about player 2's strategy choice as the one in Example 2.1, namely, L. Since player 2 has perfect knowledge about the situation, she can deduce from her information player 1's strategy choice, which is identical to the one in Example 2.1. So, if player 2 is asked to first choose between the game in Example 2.1 and the one in Example 2.3 and then to play the chosen game, there is no act-state dependence, since both players' probabilities for how the other player chooses are unchanged. For this reason player 2 would assign a positive value to the information about the state.

It has already been shown (Kadane et al., 2008) that in individual decision problems the result about the non-negative value of cost-free information does not hold provided that there is act-state dependence in personal probabilities. Then it should come as no surprise that, in the context of games, players may have negative value for new information in the presence of act-state dependence. It is act-state dependence that has driven the counterintuitive result about the negative value of information discussed above[5]. Relaxing the assumption of common knowledge in Example 2.3 enables us to prevent the occurrence of act-state dependence in the sequential problem.

3 Negative Value of Information in Games under Ambiguity

This section is devoted to extending the analysis of the value of information in Bayesian games to accommodate ambiguity aversion by considering the multiple priors models developed by Gilboa and Schmeidler (1989). In order to explore whether the phenomenon of negative value of information is robust with respect to extensions of the expected utility theory, we need to investigate how to incorporate models of imprecise probabilities into Bayesian games. We will not attempt to present a formal model of incomplete information games under ambiguity here[6]. Instead, we shall introduce the basic ideas in an informal way.

Unlike in Harsanyi (1967/68), here we take the view that, due to limited information, a player may not be able to identify a unique prior to describe her belief about the states, which can be characterized as a set of probability distributions. More precisely, we assume that in an incomplete information game each player's perception of uncertainty about the states is modeled by a closed and convex set of probability measures, instead of a single common probability distribution. And we adopt the principle of Γ-maximin as the decision rule used by all the players. In the same spirit of Bayesian Nash equilibrium, we propose a new solution concept in which each player chooses the optimal action in the sense of maximizing the minimum expected utility for each realization of her private signal. In this sense, our model constitutes only a minor departure from the standard approach to Bayesian games.

In order to help better understand how our game model works, let us consider the following example through which we introduce the major components of an incomplete information game under ambiguity and the solution concept for this kind of games.

Example 3.1. *Consider the game structure shown in Figure 3 (Game 4). As in Bayesian games, nature moves first and chooses which state to occur. Moreover, assume that both players in this game are uncertain about nature's*

[5]The author thanks Teddy Seidenfeld for pointing this out. Seidenfeld (2009) first gives such an analysis of Osborne's example regarding negative value of information in games.

[6]See Kajii and Ui (2005) for a formal definition of such a game model.

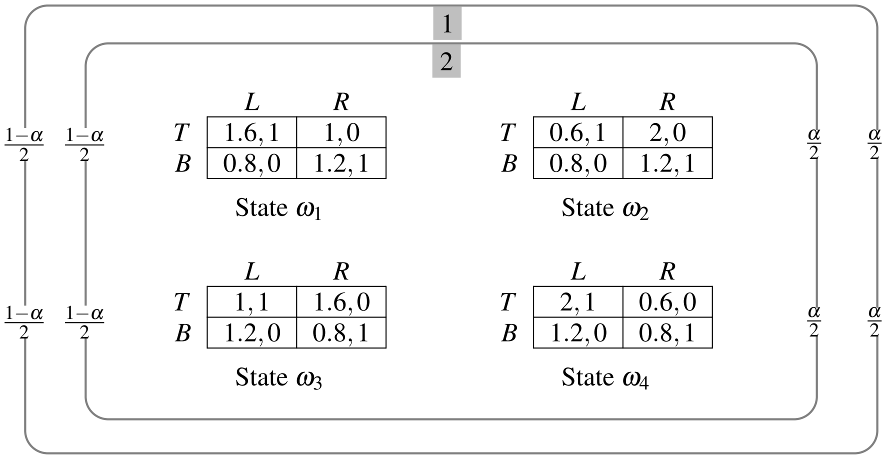

Figure 3: An incomplete information game under ambiguity

move in choosing the state. Nevertheless, suppose that both players' prior beliefs about the states are depicted by the following common set of priors over the states with $\alpha \in [0.1, 0.9]$.

$$\mathscr{P} = \{p \in \Delta(\Omega) : p(\omega_1) = p(\omega_3) = \frac{1-\alpha}{2}, \quad p(\omega_2) = p(\omega_4) = \frac{\alpha}{2}\}. \tag{1}$$

Like in the framework of Bayesian games, the solution concept proposed here requires that each player's conjecture regarding the opponents' choices is correct in the standard sense. Roughly speaking, our solution concept called Γ-*maximin equilibrium* is defined as a strategy profile such that each player's strategy is optimal in the sense of maximizing the minimum expected payoff, given the other players' strategy choices. Clearly, the concept of Γ-maximin equilibrium generalizes the notion of Bayesian Nash equilibrium to games under ambiguity.

Now let us apply the concept of Γ-maximin equilibrium to the game in Figure 3. We can transform the game into an ordinary game in strategic form by calculating each player's expected payoffs under different strategy profiles using each of her posterior probabilities defined over the state. Given a strategy profile, each player's payoff generally becomes an interval instead of a precise value. However, it turns out that in this case the players' payoffs are all determinate and independent of the variable α. By a simple calculation, it follows that this game can be turned into the following 2×2 game in strategic form.

	L	R
T	1.3, 1	1.3, 0
B	1, 0	1, 1

Figure 4: Game 4 in strategic form

It is obvious that this game can be easily solved by strict dominance, which leads to a unique solution, namely, the strategy profile (T, L). Hence, this is the unique Γ-maximin equilibrium for the incomplete information game under ambiguity in Figure 3, where its corresponding outcome is $(1.3, 1)$. One may notice that the principle of Γ-maximin does not play a role in solving this specific game. It is thus worthwhile to point out that this is not generally true for incomplete information games under ambiguity, since the payoffs would typically form intervals. As we shall see in the next example, the game is solved by explicitly applying the idea of Γ-maximin.

It has already been demonstrated that in single-agent decision problems a non-expected utility decision maker, especially a Γ-maximin decision maker, may prefer less information to more[7]. Thus one should expect that a similar phenomenon would arise in incomplete information games under ambiguity. In the following, we present a case where a Γ-maximin player would rationally pay not to receive cost-free information, which is exactly in the same spirit as the negative result presented in the previous section.

In order to draw the needed contrast with the game in Figure 3, we consider a situation in which player 2 has more information and player 1 knows that fact. In this sense, we follow exactly the same construction as Osborne has proposed for the Bayesian case.

Example 3.2. *Consider the game structure shown in Figure 5, which is similar to the game in Example 3.1. Likewise, assume that both players' prior beliefs about the states are represented by the same set of priors over the states given in Equation* (1). *As opposed to the previous case, we assume in this game that player 1 learns more information about the state, whereas player 2 does not. To be more specific, player 1 may receive two signals; when player 1*

[7] See Wakker (1988), Seidenfeld (2004) and AI-Najjar and Weinstein (2009) for various examples that illustrate this point.

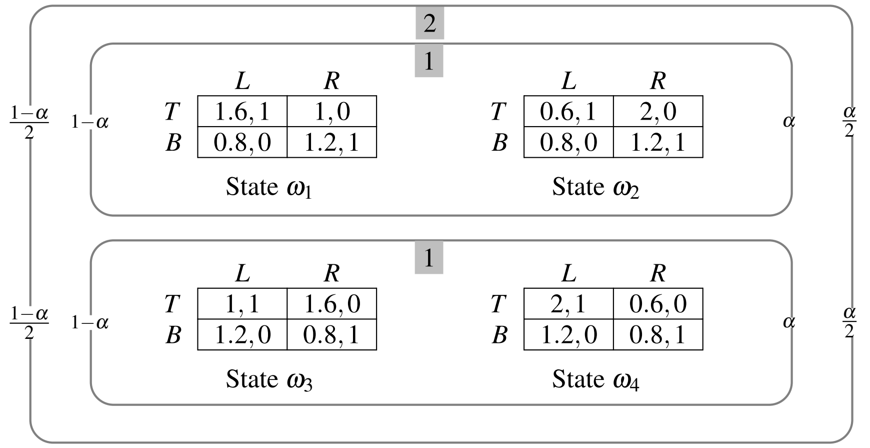

Figure 5: The second incomplete information game under ambiguity

gets one of the signals, she knows that the state is either ω_1 *or* ω_2*; when she gets another signal, she knows that the state is either* ω_3 *or* ω_4*. Formally, player 1's signal function can be defined as follows* $\tau_1(\omega_1) = \tau_1(\omega_2) = t_1$ *and* $\tau_1(\omega_3) = \tau_1(\omega_4) = t_1'$*. By contrast, player 2's signal function is given as follows:* $\tau_2(\omega_k) = t_2$ *for* $k = 1,2,3,4$*. This indicates that player 2 receives a single signal in the interim stage.*

Similarly, we want to solve this incomplete information game under ambiguity by considering the concept of Γ-maximin equilibrium. However, it is important to remark that, as opposed to the notion of Bayesian Nash equilibrium, this solution concept is sensitive to whether we solve the game from an *ex ante* or *interim* perspective. It is important to note that the ex ante and interim Γ-maximin equilibrium may lead to rather different solutions to the same incomplete information game under ambiguity, since the normal form and extensive form of a decision problem are not equivalent under Γ-maximin[8]. Here we shall only focus on the concept of interim Γ-maximin equilibrium, which can be regarded as a direct generalization of interim Bayesian Nash equilibrium.

Before introducing this interim solution, we need to specify how the players would update their beliefs upon receiving new information, which is critical for determining players' expected payoffs in the interim stage. The updating problem lies at the heart of the theory of incomplete information games. In the framework of Bayesian games, it is widely accepted that Bayes' rule provides a convenient and useful way of revising players' initial beliefs in the light of new information. On the contrary, there has been little agreement in the literature on how to update one's beliefs in the presence of ambiguity as new information is gathered[9]. Nevertheless, the so-called *full Bayesian updating rule* is often regarded as the straightforward generalization of Bayes' rule to the context of imprecise probabilities. Thus we will apply this rule to incomplete information games with the modeling of uncertainty through sets of probabilities.

Given player 1's private information, the set of posterior probabilities can be derived from the set of priors by applying the full Bayesian updating rule:

$$\mathscr{P}(\cdot|t_1) = \{p \in \Delta(\Omega) : p(\omega_1) = 1-\alpha, p(\omega_2) = \alpha\}$$
$$\mathscr{P}(\cdot|t_1') = \{p \in \Delta(\Omega) : p(\omega_3) = 1-\alpha, p(\omega_4) = \alpha\},$$

where $\alpha \in [0.1, 0.9]$. And it is clear that player 2's initial beliefs about the state would remain unchanged. Moreover, note that player 1's payoffs to some strategy profiles depends upon the value of α and thus become intervals. This stands in direct contrast with the game in Figure 3.

Similarly to the notion of interim Bayesian Nash equilibrium, the concept of interim Γ-maximin equilibrium requires that each type of each player chooses a strategy that maximizes her minimum *interim expected payoff* given the strategies chosen by the other types of every other player. We claim that this game has a unique interim Γ-maximin equilibrium in pure strategy, namely, the strategy profile $((B,B),R)$[10]. To see this, note that for each type of player 1 the action B always yields a higher minimum expectation than T does, no matter whether player 2 chooses L or R. Thus, player 1 should choose to play the strategy (B,B). Given such a conjecture, player 2 should respond optimally by playing R. Importantly, observe that this unique equilibrium yields type t_1 of player 1 an expected payoff of 1.2

[8]Seidenfeld (1988) explicitly shows that the rule of Γ-maximin distinguishes between sequential decisions in extensive form and the normal form one-stage decisions.

[9]See Gilboa and Schmeidler (1993) and Grove and Halper (1998) for more detailed discussions about various updating rules when using sets of probabilities to model ambiguity.

[10]In fact, this game has no other interim Γ-maximin equilibrium in mixed strategy. For the current purpose, however, the notion of interim Γ-maximin equilibrium in pure strategy is sufficient. Moreover, one can easily verify that this equilibrium cannot be justified a Bayesian Nash equilibrium using any element of the set of priors.

and type t_1' of player 1 an expected payoff of 0.8, both of which are less than the payoff in the unique Γ-maximin equilibrium of the game in Figure 3.

Now if player 1 has an initial choice between the game in Example 3.1 and the game in Example 3.2, player 1 strictly prefers to play the former one in which player 1 has no information about the state. In other words, player 1 assigns a negative value to the information that she may receive in the latter game. Unsurprisingly, then, this example demonstrates that the phenomenon of the negative value of information would arise in the context of incomplete information games under ambiguity. In this sense, such a phenomenon is indeed robust with respect to extensions of Bayesian games using certain classes of non-expected utility models.

In view of the analysis in Section 2, one may think that this is due to the same fact that act-state dependence obtains in this sequential problem. This is correct to some extent, since player 1's new information about the state does have a strategic impact on her conjecture about player 2's strategy choice. As we shall see in the next section, however, act-state dependence is not the key factor that accounts for this negative result in games under ambiguity. The use of imprecise probabilities to represent uncertainty in games actually introduces additional complexity to the value of information in games.

4 Negative Value of Information without Common Knowledge

In this section we consider a variant of the game in Example 3.2 where the assumption of common knowledge is slightly weakened, and demonstrate that the conclusion of the previous section still holds even if some player has more information that is not commonly known to the other players. Then, on the basis of previously known results from decision theory, we provide a more in-depth account of the nature of the negative value of information in incomplete information games with and without ambiguity.

As we announced, we reexamine the information value problem in games by comparing the game in Example 3.1 with a game where player 1 has more information but player 2 is not informed of this fact. Similarly, the construction of the latter game is deliberately designed to eliminate the strategic effects caused by common knowledge of new information.

Example 4.1. *Consider a slightly modified version of the incomplete information game under ambiguity introduced in Example 3.2. The modified game is very similar to the one in Figure 5 except that player 2's belief about player 1's information in the interim stage does not match up with the actual information possessed by player 1. More precisely, we assume here that player 1 obtains more information about the state, but that is not revealed to player 2. Instead, player 2 still thinks that player 1 holds the same belief as in the ex ante stage.*

It is worth emphasizing that the fundamental difference between Example 3.2 and the current example lies in the question of whether or not player 2 knows the fact that player 1 has more information about the state in the interim stage. As we have seen, such a modification has a profound impact on the solution to the games, which in turn substantially changes the analysis of the information value in games.

In a similar fashion, we propose to solve this game by demanding only that the strategy chosen by each player is optimal in the sense of maximizing the lower expectation on the basis of her information about the state and the other players' strategy choices. In contrast with Γ-maximin equilibrium, this notion does not impose the consistency requirement on each player's conjectures about the other players' strategy choices, since some player may not have all the information about the opponents' characteristics. For instance, in the current example, the fact that player 1 learns more information about the state is not available to player 2. For this reason it is not surprising to see that some player may hold a belief that is inconsistent with the opponents' actual behavior. In light of such a weakening of the common knowledge assumption, however, it seems quite reasonable to require each player to justify the strategy choice on the grounds of information that is available to her.

We first argue that player 2 should choose to play L. The argument goes like this. In this case, player 2 still believes that neither player obtains any further information about the state, and both of them employ the same set of priors $\mathscr{P}$ to represent their uncertainty about the state. Given such a belief, player 2 will then think that the game can be transformed into the strategic form game depicted in Figure 4, which is solvable using strict dominance. Player 2 would thus expect player 1 to select T and then choose to play her unique best response L.

Second, we claim that player 1's optimal strategy is (B,B). Unlike player 2, player 1 not only has more information about the state in the interim stage, but also knows that player 2 does not learn of this fact. Based on her private information, player 1 can reason as follows. As noted before, the action B always gives player 1 a higher minimum expected payoff than T does, regardless of the action chosen by player 2. Thus, the optimal choice for player 1 is to choose the strategy (B,B). We should also remark that, according to Γ-maximin, player 1's strategy (B,B) is optimal as well given player 2's choice of L. This means that player 1 does not want to alter her choice, even though she can expect that player 2 will choose to play L instead of the equilibrium strategy R. Importantly, note that under the specified optimal plays for both players, the payoffs to two types of player 1 are 0.8 and 1.2 respectively.

We now turn our attention to the question of whether player 1 prefers more information to less without the common knowledge assumption. Suppose that player is presented with the following sequential problem: First, decide whether to play the game in Example 3.1 or instead to play the game in Example 4.1, and then enter the selected game. What would be player 1's initial choice? Our foregoing analysis of these two games suggests that the initial choice for player 1 is to play the game in Figure 3 in which player 1 has no information about the state. That is, player 1 would rationally pay to avoid learning more information about the state in this case. Thus, the same phenomenon occurs again in the presence of ambiguity, even if player 2 is assumed to not know the fact that player 1 has more information. In contrast to the Bayesian case, one cannot explain away this counterintuitive result by weakening the assumption of common knowledge.

Let us reflect on this negative result obtained in games under ambiguity. First, it is important to note that the relaxation of the common knowledge assumption does block the kind of strategic effects caused by new information. For instance, player 2's strategy choice in the game of Example 4.1 is not affected by player 1's extra information about the state, since player 2 does not learn that fact. In this respect, the current case is quite similar to the Bayesian game. It thus follows that act-state dependence does disappear when we restrict player 2 from knowing that player 1 has more information. In the absence of act-state dependence, we still infer that player 1 strictly prefers less information to more when ambiguity is present. Taken these together, it suggests that there must be something other than act-state dependence, which leads to the undesirable result of the negative value of information in games under ambiguity.

A natural question then arises: What is the real reason behind this counterintuitive result? To address this question, we first point out a distinctive feature of the game introduced in Example 3.2 by comparing the two sets of probabilities that are used to represent player 1's prior and posterior beliefs about the state. At the beginning of the game, player 1's belief about the state is represented by the set $\mathscr{P}$ as depicted by Equation (1). By contrast, upon arrival of new information player 1 changes the probabilities by applying the full Bayesian updating rule, which gives rise to the sets $\mathscr{P}(\cdot\,|\,t_1)$ and $\mathscr{P}(\cdot\,|\,t_1')$. Observe that the prior probability interval for event $\{\omega_1,\omega_4\}$ is strictly contained within its conditional probability interval given by $\mathscr{P}(\cdot\,|\,t_1)$. Intuitively, this means that player 1's probability judgment about the event $\{\omega_1,\omega_4\}$ becomes more imprecise after she observes t_1. We can say the same thing about the probability interval for $\{\omega_2,\omega_3\}$. In a similar fashion, the same facts can be reported regarding player 1's probability estimates for these events in the case that t_1' is observed. We can thus conclude that player 1's probability intervals for nature's choice of the state become wider, regardless of the signal revealed to her. Such a phenomenon is called *dilation of sets of probabilities*.

Intuitively, one may expect that one's probability interval for some hypothesis should become narrower after learning the outcome of some experiment. Contrary to common sense, when dilation obtains, the probability interval actually expands, no matter what the outcome is. It is not surprising to find that a Γ-maximin decision maker may refuse to learn new information when dilation is present. Because the agent becomes more uncertain, no matter what the outcome of the experiment is. In single-agent decision problems, Γ-maximin requires that a decision maker using that decision rule should pay to not receive new information whenever dilation occurs.

In the context of games, it is reasonable to expect that this would also arise in the presence of ambiguity if dilation occurs. As explained above, this is exactly what happens in the example of incomplete information games under ambiguity. Hence, it is the phenomenon of dilation, together with the use of Γ-maximin, rather than act-state dependence that has really driven the result of the negative value of information in games under ambiguity. In addition, it is important to remark that in the presence of dilation the same result holds with or without the common knowledge assumption. This is why we arrive at the same conclusion in both comparisons that player 1 would rather choose the situation in which she does not have more information. As our examples have shown, when considering whether new information has any value in games, act-state dependence is very sensitive to strategic impact whereas dilation does not. In this sense, dilation is more robust than act-state dependence regarding the result about the negative value of information in the context of games under ambiguity.

5 Summary

The discussion here is concerned with the issue of whether the familiar result about the non-negative value of information is still valid under strategic situations. Osborne (2004) has shown that in Bayesian games more information to one player may make her worse off in terms of equilibrium payoffs. This finding raises the concern that a wide variety of game-theoretic models with important economic applications may fail to satisfy a seemingly reasonable requirement on rational choice. In this paper we have reexamined the information value problem in the context of incomplete information games by considering the strategic effect of common knowledge on players' strategies and also the introduction of ambiguity.

We draw two conclusions from this investigation. First, the role act-state dependence plays in the sequential setting is fundamental to understanding the result that a Bayesian player may rationally pay not to receive cost-free information in strategic interactions. We have argued that the kind of probabilistic dependence between a player's initial choice of the games and her probability about the opponents' strategy choice leads the player to assign a negative value to the information. We have further shown that the non-negative value of information can be restored by weakening the assumption of common knowledge. This is mainly due to the fact that the lack of common knowledge isolates the strategic effect of information on equilibrium play, and thus removes act-state dependence. So this result lends an additional support for our account of the result of the negative value of information in Bayesian games.

Second, in the presence of ambiguity, more information may also damage the player who holds it, thereby implying a negative value of information. Yet, we should emphasize that the negative value of information can occur in the context of incomplete information games under ambiguity, even if we isolate the effect of act-state dependence by relaxing the common knowledge assumption. Unlike what happens in Bayesian games, in this case the result about the negative value of information still holds true mainly due to dilation of sets of probabilities. The upshot is that, within the generalized game-theoretic framework that allows for the modeling of uncertainty through sets of probabilities, one needs to pay attention to both the strategic effect of common knowledge and the role of ambiguity in order to respect the value of information.

Acknowledgements

I would like to thank Teddy Seidenfeld and Kevin Zollman who originally suggested this topic, and provided constant encouragement and valuable suggestions throughout the investigation. I am grateful to Martin Osborne for helping me with a technical problem when using the latex package sgame.sty created by him. I also want to thank the four anonymous reviewers for their insightful comments on an earlier version of the paper.

References

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84(3), 488-500.

Al-Najjar, N. I., and Weinstein J. (2009). The Ambiguity Aversion Literature: A Critical Assessment. *Economics and Philosophy* 25, 249-284.

Ellbsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics* 75, 643-649.

Gilboa, I., and Schmeidler, D. (1989). Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics* 18, 141-153.

Gilboa, I., and Schmeidler, D. (1993). Updating Ambiguous Beliefs. *Journal of Economic Theory* 59, 33-49.

Good, I. J. (1967). On the Principle of Total Evidence. *British Journal of Philosophy of Science* 17, 319-321.

Grove, A. J., and Halpern, J. (1998). Updating Sets of Probabilities. In *Proceedings of the Fourteenth Conference on Uncertainty in AI*, 173-182.

Harsanyi, J. C. (1967/68). Games with Incomplete Information Played by 'Bayesian' Players. *Management Science* 14, 159–182, 320–334, and 486–502.

Herron, T., Seidenfeld, T., and Wasserman, L. (1994). The Extent of Dilation of Sets of Probabilities and the Asymptotics of Robust Bayesian Inference. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, Vol. 1, 250-259.

Herron, T., Seidenfeld, T., and Wasserman, L. (1997). Divisive Conditioning: Further Results on Dilation. *Philosophy of Science* 64, 411-444.

Kadane, J. B., Schervish, M., and Seidenfeld, T. (2008). Is Ignorance Bliss? *The Journal of Philosophy* 105, 5-36.

Kajii, A., and Ui, T. (2005). Incomplete Information Games with Multiple Priors. *Japanese Economic Review* 56, 332-351.

Knight, F. H. (1921). *Risk, Uncertainty, and Profit.* Boston, MA: Houghton Mifflin Company.

Levi, I. (1974). On Indeterminate Probabilities. *Journal of Philosophy* 71, 391-418.

Neyman, A. (1991). The Positive Value of Information. *Games and Economic Behavior* 3, 350-355.

Osborne, M. J. (2004). *An Introduction to Game Theory*. New York, NY: Oxford University Press.

Ramsey, F. P. (1990). Weight or the Value of Knowledge. *British Journal for the Philosophy of Science* 41, 1-4.

Seidenfeld, T. (1988). Decision Theory without "Independence" or without "Ordering": What is the Difference? *Economics and Philosophy* 4, 267-290.

Seidenfeld, T. (2004). A Contrast between Two Decision Rules for Use with (Convex) Sets of Probabilities: Γ-maximin versus E-admissibility. *Synthese* 140, 69-88.

Seidenfeld, T. (2009). On the Value of Information in Games (Online). Available: www.hss.cmu.edu/philosophy/faculty-seidenfeld.php.

Seidenfeld, T., and Wasserman, L. (1993). Dilation for Sets of Probabilities. *Annals of Statistics* 21, 1139-1154.

Wakker, P. (1988). Non-Expected Utility as Aversion to Information. *Journal of Behav Decis Mak* 1, 169-175.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. New York: Chapman and Hall.